

# Computational Physics on Graphics Processing Units

Ari Harju<sup>1,2</sup>, Topi Siro<sup>1,2</sup>, Filippo Federici Canova<sup>3</sup>, Samuli Hakala<sup>1</sup>, and Teemu Rantalaiho<sup>2,4</sup>

<sup>1</sup> COMP Centre of Excellence, Department of Applied Physics, Aalto University School of Science, Helsinki, Finland [ari.harju@aalto.fi](mailto:ari.harju@aalto.fi)

<sup>2</sup> Helsinki Institute of Physics, Helsinki, Finland

<sup>3</sup> Department of Physics, Tampere University of Technology, Tampere, Finland

<sup>4</sup> Department of Physics, University of Helsinki, Helsinki, Finland.

Submitted to Lecture Notes in Computer Science.

We acknowledge comments, suggestions, and missing references to

[ari.harju@aalto.fi](mailto:ari.harju@aalto.fi)

**Abstract.** The use of graphics processing units for scientific computations is an emerging strategy that can significantly speed up various different algorithms. In this review, we discuss advances made in the field of computational physics, focusing on classical molecular dynamics, and on quantum simulations for electronic structure calculations using the density functional theory, wave function techniques, and quantum field theory.


**Keywords:** graphics processing units, computational physics

## 1 Introduction

The graphics processing unit (GPU) has been an essential part of personal computers for decades. Their role became much more important in the 90's when the era of 3D graphics in gaming started. One of the hallmarks of this is the ultimately violent first-person shooting game called DOOM by the id Software company, released in 1993. Wandering around the halls of slaughter, it was hard to imagine these games leading to any respectable science. However, now twenty years after the release of DOOM, the gaming industry has an enormous market, and the continuous need for more realistic visualizations has lead to a situation where the GPUs have, indeed, an enormous computational power, much higher than the central processing units (CPUs), at least in terms of theoretical peak performance.

The games started to have real 3D models and hardware acceleration in the mid 90's, but an important turning point for scientific use of GPU's for computing was around the first years of this millennium [1], when the programmability

of GPU's was introduced. Combined with the continued increase in computational power as shown in Fig. 1, the GPUs are nowadays a serious platform for general purpose computing. Also, the memory bandwidth in GPUs is very impressive. The three main vendors for GPUs, Intel, NVIDIA, and ATI/AMD, are all actively developing computing on GPUs. In addition, the Cell processor architecture by IBM, Sony, and Toshiba is also relevant for scientific activities and often grouped with the GPU efforts. At the moment, none of the technologies listed above dominate the field, but NVIDIA with its CUDA programming environment is perhaps the market leader at the moment.



**Fig. 1.** Floating point operations (FLOPS) per second for GPUs and CPUs from NVIDIA and Intel Corporations, figure taken from[2]. The processing power of the currently best GPU hardware by the AMD Corporation is comparable to the NVIDIA, being around 2600 GFLOPS/s.

### 1.1 GPU as a computational platform

At this point we have hopefully convinced you that the GPU is a powerful architecture also for general computing, but what makes GPU different from the current multi-core CPUs? To understand this, we can start with traditional graphics processing, where hardware vendors have tried to maximize the speed at which the pixels on the screen are calculated. These pixels are independent primitives that can be processed in parallel, and the number of pixels of computer

displays has also increased over the years from the original DOOM resolution of  $320 \times 200$  corresponding to 64000 pixels to millions. The most efficient way to process these primitives is to have a very large number of arithmetic logical units (ALUs) that are able to perform a high number of operations for each video frame. The processing is very data-parallel, and one can view this so that the same arithmetic operation is performed in parallel for each primitive to be processed. Furthermore, as the operation is the same for each primitive, there is no need for very sophisticated flow control in GPU and more transistors can be used for arithmetics, resulting in an enormously efficient hardware for performing parallel computing that can be classified as “single instruction, multiple data” (SIMD).

Now, for general computing on the GPU, the primitives processed are not anymore the pixels on the video stream, but can range from matrix elements in linear algebra to physics related cases where the primitives can be particle coordinates in classical molecular dynamics or quantum field values. From the traditional graphics processing we can understand that the general computing would be efficient when we have a situation where the same arithmetics needs to be performed for a large data set. It is clear that not all problems or algorithms have this structure, but there are luckily many cases where this is the case, and the list of successful examples is long.

However, there are also limitations on GPU computing. First of all, when porting a CPU solution for a given problem to GPU, one might need to change the algorithm to suit the SIMD approach. Secondly, the communication from the host part of the computer to the GPU part is limited by the speed of the PCIe bus coupling the GPU and the host. In practice, this means that one needs to perform a serious amount of computing on the GPU between the data transfers before the GPU can actually speed up the overall computation. Of course, there are also cases where the computation as a whole is done on GPU, but these cases suffer from the somewhat smaller serial processing power of the GPU.


## 2 Molecular dynamics

Molecular dynamics (MD) refers to the type of simulation where the behaviour of a complex system is calculated by integrating the equation of motion of its components within any given model, and its goal is to observe how some ensemble-averaged properties of the system originate from the detailed configuration of its constituent particles (figure 2).

In its classical formulation, the dynamics of a system of particles is described by their Newtonian equations:

$$m_i \frac{d^2 \mathbf{x}_i}{dt^2} = \sum_j \mathbf{F}_{ij} \quad (1)$$

where  $m_i$  is the particle’s mass,  $\mathbf{x}_i$  its position, and  $\mathbf{F}_{ij}$  is the interaction between the  $i$ -th and  $j$ -th particles as provided by the model chosen for the studied



**Fig. 2.** Schematic representation of the atomistic model of a macroscopic system.

system. These second order differential equations are then discretised in the time domain, and integrated step by step until one is satisfied with the results.

The principles behind MD are so simple and general that since its first appearance in the 70s, it has been applied to a wide range of systems, at very different scales. Often MD is applied in the field of biophysics, since it happens to be the dominant theoretical tool of investigation, where structural changes in proteins [3,4,5,6] and lipid bilayers [7,8] interacting with drugs can be studied, ultimately providing a better understanding of drug delivery mechanisms.

At larger scales, one of the most famous example is known as the *Millenium Simulation*, where the dynamics of the mass distribution of the universe at the age of 380000 years, was simulated up to the present day [9], showing the formation age and process of cosmic structures (galaxies, black holes and quasars), greatly improving our understanding of cosmological models and providing a theoretical comparison to satellite measurements.

Despite the simplicity and elegance of its formulation, MD is not a computationally easy task and often requires special infrastructure. The main issue is usually the evaluation of all the interactions  $\mathbf{F}_{ij}$  that is the most time consuming procedure of any MD calculation for large systems. Moreover, the processes under study might have long characteristic time scales, requiring longer simulation time and larger data storage; classical dynamics is chaotic, i.e. the outcome is affected by the initial conditions, and since these are in principle unknown and chosen at random, some particular processes of interest might not occur just because of the specific choice, and the simulation should be repeated several times. For these reasons it is important to optimise the evaluation of the forces as much as possible.

Fortunately, MD is a very suitable problem for parallel computation, as all the interactions  $\mathbf{F}_{ij}$  are independent of each other and their evaluation can be distributed over different computing units; custom hardware architectures dedicated to a specific MD simulation were also developed [10,11,12], and despite the good performance achieved, their excessive cost and lack of general applica-


bility hampered their spread through scientific communities, or even beyond the particular scope they were initially intended for.

Recently, GPUs attracted a lot of interest from the scientific community because they feature a massively parallel architecture with performance on the level of small computer clusters, at the cost and power consumption of conventional commodity hardware. Even though they were originally designed for real-time 3D rendering, the computer games market drove GPU manufacturers to fit ever more processing power on their devices, to the point where it became advantageous to calculate the game physics (dynamics and collisions) on the device itself [13,14].

An early attempt to implement MD on GPU was proposed in 2004 [15] and showed promising performance; at that time, general purpose GPU computing was not yet a well established framework and the N-body problem had to be formulated as a rendering task: a *shader* program computed each pair interaction  $\mathbf{F}_{ij}$  and stored it as the pixel color values (RGB) in a  $N \times N$  texture, then another shader would simply sum these values row-wise to obtain the total force on each particle and finally integrate their velocities and positions. The method is called all-pairs calculation, and as the name might suggest it is quite expensive as it requires  $\mathcal{O}(N^2)$  forces evaluations. The proposed implementation was in no way optimal since the measured performance was about a tenth of the nominal value of the device, and immediately revealed one of the main issues of the architecture that still persists nowadays: GPUs are very fast at computing data, but extremely slow at handling them. The reason for the bad performance was in fact the large amount of memory read instructions compared to the amount of computation effectively performed on the fetched data, but despite this limitation, the code still outperformed a CPU by a factor of 8 because every interaction was computed concurrently. A wide overview of optimisation strategies to get around the memory latency issues can be found in Ref. [16], while, for the less eager to get their hands dirty, a review of available MD software packages is included in Ref. [17].

In the current GPU programming model, the computation is distributed in different threads, grouped together as blocks in a grid fashion, and they are allowed to share data and synchronise throughout the same block; the hardware also offers one or two levels of cache to enhance data reuse thus reducing the amount of memory accesses, without harassing the programmer with manual pre-fetching. A more recent implementation of the all-pair calculation [18] exploiting the full power of the GPU can achieve performance close to the nominal values, comparable to several CPU nodes.

The present, more mature, GPGPU framework allows for more elaborate programs to fit the device, enabling the implementation of computational *tricks* developed during the early days of MD [19], in order to integrate N-body dynamics accurately with much better scaling than  $\mathcal{O}(N^2)$ . For example, in many cases the inter-particle forces are short range, and it would be unnecessary to evaluate every single interaction  $\mathbf{F}_{ij}$  since quite many of them would be close to zero and just be neglected. It is good practice to build lists of neighbours for



**Fig. 3.** Illustration of different space partition methods. In dense systems (a) a regular grid is preferred, and neighbour particles can be searched only in a few adjacent voxels. Sparse system (b) are better described by hierarchical trees, excluding empty regions from the computation.

each particle in order to speed up the calculation of forces: this also takes an  $\mathcal{O}(N^2)$  operation, although the list is usually recalculated every 100-1000 time-steps, depending on the average mobility of the particles. Another way to exclude distant pairs from the calculation is to divide the simulation box in voxels and search for a particle's neighbours only within the adjacent voxels (figure 3a). Performance can be further improved by sorting particles depending on the index of the voxel they belong, making neighbouring particles in space, to a degree, close in memory thus increasing coalescence and cache hit rate; such a task can be done with radix count sort [20,21,22] in  $\mathcal{O}(N)$  with excellent performance, and it was shown to be the winning strategy [23].

Unfortunately, most often the inter-particle interactions are not exclusively short range and can be significant even at larger distances (electrostatic and gravitational forces). Therefore introducing an interaction cut-off leads to the wrong dynamics. For dense systems, such as bulk crystals or liquids, the electrostatic interaction length largely exceeds the size of the simulation space, and in principle, one would have to include the contributions from several periodic images of the system, although their sum is not always convergent. The preferred approach consists of calculating the electrostatic potential  $V(\mathbf{r})$  generated by the distribution of point charges  $\rho(\mathbf{r})$  from Poisson's equation:

$$\nabla^2 V(\mathbf{r}) = \rho(\mathbf{r}) \quad (2)$$

The electrostatic potential can be calculated by discretising the charge distribution on a grid, and solve equation 2 with fast Fourier transform (FFT), which has  $\mathcal{O}(M \log M)$  complexity (where  $M$  is the amount of grid points): this approach is called particle-mesh Ewald (PME). Despite being heavily non-local, much work has been done to improve the FFT algorithm and make it cache efficient [24,25,26,27,28], so it is possible to achieve a 20-fold speed up over the

standard CPU FFTW or a 5-fold speedup when compared to highly optimised MKL implementation. The more recent multilevel summation method (MSM) [29] uses nested interpolations of progressive smoothing of the electrostatic potential on lattices with different resolutions, offering a good approximation of the electrostatic  $\mathcal{O}(N^2)$  problem in just  $\mathcal{O}(N)$  operations. The advantage of this approach is the simplicity of its parallel implementation, requiring less memory communication among the nodes, thus leading to a better scaling than FFT calculation in PME. The GPU implementation of this method gave a 25-fold speedup over the single CPU [30]. Equation 2 can also be translated into a linear algebra problem using finite differences, and solved iteratively on multi-grids [31,32] in theoretically  $\mathcal{O}(M)$  operations: even though the method initially requires several iterations to converge, the solution does not change much in one MD step and can be used as a starting point in the following step, which in turn will take much fewer iterations.

On the other hand, for sparse systems such as stars in cosmological simulations, decomposing the computational domain in regular boxes can be quite harmful because most of the voxels will be empty and some computing power and memory is wasted there. The optimal way to deal with such situation is to subdivide space hierarchically with an octree [33] (figure 3b), where only the subregions containing particles are further divided and stored. Finding neighbour particles can be done via traversal of the tree in  $\mathcal{O}(N \log N)$  operations. Octrees are conventionally implemented on the CPU as dynamical data structures where every node contains reference pointers to its parent and children, and possibly information regarding its position and content, although this method is not particularly GPU friendly since data is scattered in memory as well as in the simulation space. In GPU implementations, the non-empty nodes are stored as consecutive elements in an array or texture, and include the indices of children nodes [34] and proved to give good acceleration in solving the N-body problem [18,35,36]. Long range interactions are then calculated explicitly for the near neighbours, while the fast multipole method (FMM) [37,38] can be used to evaluate contributions from distant particles. The advantage of representing the system with an octree becomes now more evident: there exists a tree node containing a collection of distant particles, which can be treated as a single multipole leading to an overall complexity  $\mathcal{O}(N)$ . The mathematics required by FMM is quite intensive to evaluate, nevertheless the algorithms involved have been developed and extensively optimised for GPU architecture [39,40,41], achieving excellent parallel performance even on large clusters [42].

In all the examples shown here, the GPU implementation of a method outperformed its CPU counterpart: in many cases the speedup is only 4-5 fold when compared to a highly optimised CPU code, which is, in a way, a discouraging result. This is so because implementing an efficient GPU algorithm is quite a difficult task, requiring knowledge of the target hardware, and the programming model is not as intuitive as regular CPU: the implementation difficulty could be eased by developing better compilers, that check how memory is effectively accessed and provide higher levels of GPU optimisation on older CPU codes

automatically. In some cases 10 to 100 fold speedup was measured, showing that GPU is clearly the way to go for a defined group of problems, while it fails for others. This issue is quite similar to the everlasting dispute between raytracing and raster graphics: one can calculate explicitly photorealistic images in complex scenes, taking its time (CPU), while the latter resorts to every trick in the book to get a visually "alright" result as fast as possible (GPU). Best would be to use both methods to calculate what they are good for, and this sets a clear view of the future hardware required for scientific computing, where both simple vector-like processors and larger CPU cores could access the same memory resources, avoiding data transfer.

### 3 Density-functional theory

Density functional theory (DFT) is a popular method for ab initio electronic structure calculations in material physics and quantum chemistry. We use the following density-functional theory formulation for the ground state Kohn-Sham wave-functions  $\psi_n(\mathbf{r})$ :

$$H\psi_n(\mathbf{r}) = \epsilon_n\psi_n(\mathbf{r}), \quad (3)$$

where  $H = -\frac{1}{2}\nabla^2 + v_H(\mathbf{r}) + \mathbf{v}_{\text{ext}}(\mathbf{r}) + \mathbf{v}_{\text{xc}}(\mathbf{r})$  and  $v_H$  is the Hartree potential defined by the Poisson equation  $\nabla^2 v_H(\mathbf{r}) = -4\pi\rho(\mathbf{r})$ ,  $v_{\text{ext}}$  is the external ionic potential and  $v_{\text{xc}}$  is the exchange-correlation potential. The electronic charge density  $\rho(\mathbf{r})$  is determined by the eigenvectors  $\rho(\mathbf{r}) = \sum_i f_i |\psi_i(\mathbf{r})|^2$ , where the  $f_i$ 's are the orbital occupation numbers.

There are several numerical approaches and approximations for solving the KS equation. They relate usually to discretization of the equations and the treatment of the core electrons (pseudo-potential and all electron methods). The most common discretizations methods in solid state physics are plane waves, localized orbitals, real space grids and finite elements. Generally an iterative procedure called self-consistent field (SCF) calculation is used to find the solution to the eigenproblem starting from initial guess for the charge density.

Porting existing DFT code to GPUs generally include profiling or discovering with some other method the most computationally expensive parts of SCF-loop and reimplementing them with GPUs. Depending on the discretization methods the known numerical bottlenecks are vector operations, matrix products, Fast Fourier Transforms (FFTs) and stencil operations. There are GPU versions of many of the standard computational libraries (like CUBLAS for BLAS and CUFFT for FFTW). However generally porting DFT application is generally not as simple as replacing the function calls with GPU equivalents since the resulting intermediate data gets heavily reused.

Partial DFT implementations on GPU include calculation of the exchange-correlation term [43] and evaluation of the Coulomb potential [44] in Gaussian based systems. However the first complete DFT code using double precision arithmetic on GPUs was presented by Genovese *et al.* [45]. They used a Daubechies wavelet based code called BIGDFT [46]. The basic 3D operations for a wavelet based code are based on convolutions. They achieved speed-ups of



factor 20 for some of these operations with GPU and a factor of 6 for the whole hybrid code using NVIDIA Tesla S1070 cards. These speed-ups were achieved on a 12-node hybrid machine.

For solid state physics plane wave (PW) basis sets are most common choice. The computational schemes rely heavily on linear algebra operations and fast Fourier transformations. Vienna ab initio Simulation Package (VASP) [47] is a popular code combining plane waves with projector-augmented wave (PAW) method. The most time consuming part of optimizing the wave-functions given the trial wave-functions and related routines have been ported to GPUs. For blocked Davinson scheme [48] and for the RMM-DIIS algorithm [49] speed-ups by a factor between 3 and 8 real-world examples with Fermi C2070 cards were achieved. Parallel scalability with 16 GPUs was similar to 16 CPUs.

Quantum ESPRESSO [50] is a electronic structure code based plane wave basis sets and pseudo-potentials (PP). For GPU version [51] the most computationally parts of SCF-cycle were gradually transferred to run on GPUs. FFTs were accelerated by CUFFT, LAPACK by MAGMA and other routines were replaced by CUDA kernels. GEMM operations were replaced by parallel hybrid phiGEMM [52] library. For single node test systems running with NVIDIA Tesla C2050 speed-ups between 5.5 and 7.8 were achieved and for 32 node parallel system the speed-ups between 2.5 and 3.5 were observed. Wand *et al.* [53] did a GPU cluster implementation of plane wave pseudo-potential code called PEtot. They were able to achieve speed-ups in computationally time by factor of 6 to 10 and parallel scalability up to 256 CPU-GPU computing units.

GPAW [54] is a density-functional theory (DFT) electronic structure program package based on grid based projector augmented wave (PAW) method. We have used GPUs to speed up most of the computationally intensive parts of solving the code: solving the Poisson equation, iterative refinement of eigenvectors, subspace diagonalization and orthonormalization of the wavefunctions. Overall we've achieved speed-ups up to 15 times on large systems and a good parallel scalability with up to 200 GPUs using NVIDIA Tesla M2070 cards.

## 4 Quantum field theory

Quantum field theories are currently our best models for fundamental interactions of the natural world (For a brief introduction to quantum field theories – or QFTs – see for example [55] or [56] and references therein). Common computational techniques includes perturbation theory, which works well in quantum field theories as long as the couplings are small enough to be considered as perturbations to free theory. Therefore, perturbation theory is the primary tool used in pure QED, weak nuclear force and high momentum-transfer QCD phenomena, but it breaks up when the coupling constant of the theory (the measure of the interaction strength) becomes large, such as in low-energy QCD.


Formulating the quantum field theory on a space-time lattice provides an opportunity to study the model non-perturbatively and use computer simulations to get results for a wide range of phenomena – it enables, for example, one to

compute the hadronic spectrum of QCD [57] from first principles and provides solutions for many vital gaps left by the perturbation theory, such as structure functions of composite particles [58], form-factors [59] and decay-constants [60]. It also enables one to study and test models for new physics, such as technicolor theories [61] and quantum field theories at finite temperature [62], [63]. For an introduction to Lattice QFT see for example [64], [65] or [66].

Simulating quantum field theories using GPUs is not a completely new idea and early adopters even used OpenGL (graphics processing library) to program the GPUs to solve lattice QCD [67] – the early GPGPU-programmers needed to set up a program that draws two triangles that fill the output-texture of desired size by running a “shader program”, which does the actual computation, for each output-pixel and in this program the input-data could be then accessed by fetching pixels input texture(s) using the texture-units of the GPU. In lattice QFT, where one typically needs to fetch the nearest neighbor lattice-site values, this actually results in good performance as the texture caches and layouts of the GPUs have been optimized for local access-patterns for filtering purposes.

#### 4.1 Solving QFTs numerically

The idea behind lattice QFT is based on the discretization of the path-integral solution to expectation values of time-ordered operators in quantum field theories. First one divides space-time into discrete boxes, called the lattice, and places the fields onto the lattice sites and onto the links between the sites, as shown in Fig. 4. Then one can simulate nature by creating a set of multiple field configurations, called an *ensemble* and calculate values of physical observables by computing ensemble averages over these states.



**Fig. 4.** The matter fields  $\Psi(x)$  live on lattice sites, whereas gauge fields  $U_\mu(x)$  live on links connecting the sites. Also depicted staples connecting to a single link variable, that are needed in computation of the gauge field forces.

The set of states is normally produced with the help of a Markov chain and in the most widely studied QFT, the lattice QCD, the chain is produced by combining a *molecular dynamics* algorithm together with a *Metropolis* acceptance test. Therefore the typical computational tasks in lattice QFTs are:

1. Refresh generalized momentum variables from a heat bath (Gaussian distribution) once per *trajectory*.
2. Compute generalized forces for fields for each step
3. Integrate classical equations of motion for the fields at each step <sup>5</sup>
4. Perform a Metropolis acceptance test at end of trajectory in order to achieve the correct limiting distribution.

In order to reach satisfying statistics, normally thousands of these trajectories need to be generated and each trajectory is typically composed of 10 to 100 steps. The force calculation normally involves a matrix inversion, where the matrix indices run over the entire lattice. It is therefore the heaviest part of the computation – the matrix arises in simulations with dynamical fermions (normal propagating matter particles) and the simplest form for the Fermion matrix is<sup>6</sup>

$$A_{x,y} = [Q^\dagger Q]_{x,y} \quad \text{where}$$

$$Q_{x,y} = \delta_{x,y} - \kappa \sum_{\mu=\pm 1}^{\pm 4} \delta_{y+\hat{\mu},x} (1 + \gamma_\mu) U_\mu(x). \quad (4)$$

Here  $\kappa$  is a constant related to the mass(es) of the quark(s),  $\delta_{x,y}$  is the *Kronecker delta function* (unit matrix elements), the sum goes over the space-time dimensions  $\mu$ ,  $\gamma_\mu$  are 4-by-4 constant matrices and  $U_\mu(x)$  are the link variable-matrices that carry the force (gluons for example), from one lattice site to the neighbouring one – in normal QCD they are 3-by-3 complex matrices.

The matrix  $A$  in the equation  $Ar = z$ , where one is solving for the vector  $r$  with given  $z$ , is an almost diagonal sparse matrix with *predefined sparsity pattern*. This fact makes lattice QCD ideal for parallelization, as the amount work done by each site is constant. The actual algorithm used in the matrix inversion is normally some variant of the conjugate gradient algorithm, and therefore one needs fast code to handle the multiplication of a fermion vector, by the fermion matrix.

This procedure is the generation of the lattice configurations which form the ensemble. Once the set of configurations  $\{U_i\}, i \in [1, N]$  has been generated with the statistical weight  $e^{-S[U_i]}$ , where  $S[U_i]$  is the *Euclidean action* (action in imaginary time formulation) expectation value of an operator  $F[U]$  can be computed simply as

$$\langle F[U] \rangle \approx \frac{1}{N} \sum_{i=1}^N F[U_i], \quad (5)$$

---

<sup>5</sup> The integration is *not* done with respect to normal time variable, but through the Markov chain index-“time”.

<sup>6</sup> There are multiple different algorithms for simulating fermions, here we present the simplest one for illustrative purposes.

## 4.2 Existing GPU Solutions to Lattice QFTs

As lattice QFTs are normally easily parallelizable they fit well into the GPU programming paradigm, which can be characterized as a parallel throughput machine. The conjugate gradient methods perform many fermion matrix vector multiplications whose arithmetic intensity (ratio of floating point operations done per byte of memory fetched) is quite low, making memory bandwidth the normal bottleneck within a single processor. Parallelization between processors is done by standard MPI domain decomposition techniques. The conventional wisdom that this helps due to higher local volume to communication surface area ratio is actually flawed, as typically the GPU can handle a larger volume in the same amount of time, hence requiring the MPI-implementation to also take care of a larger surface area in the same time as with a CPU. In our experience GPU adoption is still in some sense in its infancy, as the network implementation seems to quickly become the bottleneck in the computation and the MPI implementations of running systems seem to have been tailored to meet the needs of the CPUs of the system. Another aspect of this is that normally the GPUs are coupled with highly powered CPUs in order to cater for the situation where the users use the GPUs in just a small part of the program and need a lots of sequential performance in order to try to keep the serial part of the program up with the parallel part. The GPU also needs a lot of concurrent threads (in the order of thousands) to be filled completely with work and therefore good performance is only achievable with relatively large local lattice-sizes.

Typical implementations assign one GPU thread per site, which makes parallelization easy and gives the compiler quite a lot of room to find instruction-level parallelism, but in our experience this can result in a relatively high register pressure: The quantum fields living on the sites have many indices (normally color and Dirac indices) and are therefore vectors or matrices with up to 12 complex numbers per field per site in the case of quark fields in normal QCD. Higher parallelization can be achieved by taking advantage of the vector-like parallelism inside a single lattice-site, but this may be challenging to implement in those loops where the threads within a site have to collaborate to produce a result, especially because GPUs impose restrictions on the memory layout of the fields (consecutive threads have to read consecutive memory locations in order to reach optimal performance [2]). In a recent paper [68] the authors solve the *gauge fixing* problem using overrelaxation techniques and they report an increase in performance by using multiple threads per site, although in this case the register pressure problem is even more pronounced and the effects of register spilling to L1 were not studied.

Currently, there are various groups using GPUs to do lattice QFT simulations. Standardization efforts for high precision Lattice QCD libraries are underway and the QUDA library [69] scales to hundreds of GPUs by using a local Schwarz-preconditioning technique, effectively eliminating all the GPU-based MPI communications for a significant portion of the calculation. They employ various optimization techniques, such as *mixed-precision* solvers, where parts of the inversion process of the fermion matrix is done at lower precision of floating

point arithmetic and using reduced representations of the SU3-matrices. Scaling to multiple GPUs can also be improved algorithmically: already a simple (almost standard) *clover improvement* [70] term in the fermion action leads to better locality and of course improves the action of the model as well, taking the lattice formulation closer to the continuum limit. Domain decomposition and taking advantage of restricted additive Schwarz (RAS) preconditioning using GPUs was already studied in 2010 in [71], where the authors get best performance on a  $32^4$  lattice with vanishing overlap between the preconditioning domains and three complete RAS-iterations each containing just five iterations to solve the local system of  $4 \times 32^3$  sites. It should be noted though that the hardware they used is already old, so optimal parameters with up-to-date components could slightly differ.

Very soon after starting to work with GPUs on Lattice QFTs one notices the effects of Amdahl's law which just points out the fact that there is an upper bound for whole program performance improvement related to optimizing just a portion of the program. It is quite possible that the fermion matrix inversion takes 90% of the total compute time, but making this portion of the code run 10 times faster reveals something odd: now we are spending half of our time computing forces and doing auxiliary computations and if we optimize this portion of the code as well, we improve our performance by a factor of almost two again – therefore optimizing only the matrix inversion gives us a mere fivefold performance improvement instead of the promised order of magnitude improvement. Authors of [72] implemented practically the entire HMC-trajectory on the GPU to fight Amdahl's law and recent work [73] on the QDP++ library implements *Just-in-Time* compilation to create GPU kernels on the fly to accommodate any non-performance critical operation over the entire lattice.

Work outside of standard Lattice QCD using GPUs include implementation of the Neuberger-Dirac overlap operator [74], which provides *chiral symmetry* at the expense of a non-local action. Another group uses the Arnoldi algorithm on a multi-GPU cluster to solve the overlap operator [75] and show scaling up to 32 GPUs. Quenched SU2 [76] and later quenched SU2, SU3 and generic SU( $N_c$ ) simulations using GPUs are described in [77] and even compact U(1) Polyakov loops using GPUs is studied in [78]. Scalar field theory – the so-called  $\lambda\phi^4$ -model – using AMD GPUs is studied in [79]. TWQCD collaboration has also implemented almost the entire HMC trajectory computation with dynamical Optimal Domain Wall Fermions, which improve chiral symmetry of the action [80].

While most of the groups use exclusively NVIDIA's CUDA-implementation [2], which offers good reliability, flexibility and stability, there are also some groups using the OpenCL standard [81]. A recent study [82] showed better performance on AMD GPUs than on NVIDIA ones using OpenCL, although it should be noted that the NVIDIA GPUs were consumer variants with reduced double-precision throughput and that optimization was done for AMD GPUs. The authors of [72] have implemented both CUDA and OpenCL versions

of their staggered fermions code and they report a slightly higher performance for CUDA and for NVIDIA cards.

### 4.3 Summary

All in all, lattice QFT using GPUs is coming out of being a promising technology into a very viable alternative to traditional CPU-based computing. When reaching for the very best strong scaling performance – meaning best performance for small lattices – single threaded performance does matter if we assume that the rest of the system scales to remove other bottlenecks (communication, memory bandwidth.) In these cases it seems that currently the best performance is achievable through high-end supercomputers such as the IBM Blue Gene/Q [83], where actually the microprocessor architecture is starting to resemble more a GPU than a traditional CPU: the PowerPC A2 chip has 16 in-order cores, each supporting 4 relatively light weight threads and a crossbar on-chip network. A 17th core runs the OS functions and 18th core is a spare to improve yields or take place of a damaged core. This design gives the PowerPC A2 chip similar performance to power ratio as a NVIDIA Tesla 2090 GPU, making Blue Gene/Q computers very efficient. One of the main advantages of using GPUs (or GPU-like architectures) over traditional serial processors is the increased performance per watt and the possibility to perform simulations on commodity hardware.

## 5 Wave function methods

The stochastic techniques based on Markov chains and the Metropolis algorithm showed great success in the field theory examples above. There are also many-body wave function methods that use the wave function as the central variable and use stochastic techniques for the actual numerical work. These quantum Monte Carlo (QMC) techniques have shown to be very powerful tools to study electronic structures beyond the mean-field level of for example the density functional theory. A general overview of QMC can be found from [84]. The simplest form of the QMC algorithms is the variational QMC, where a trial wave function with free parameters is constructed and the parameters are optimized, for example, to minimize the total energy [85]. This simple strategy works rather well for various different systems, even for strongly interacting particles in an external magnetic field [86].

There has been some works porting QMC methods to GPUs. In the early work by Amos G. Anderson *et al.* [87], the overall speedup compared to CPU was rather modest, from three to six, even if the individual kernels were up to 30 times faster. More recently, Kenneth P. Esler *et al.* [88] have ported the QMC-Pack simulation code to the Nvidia CUDA GPU platform. Their full-application speedups is typically around 10 to 15 compared to a quad-core Xeon CPU. This speedup is very promising and demonstrates the great potential GPU computing has for the QMC methods that are perhaps the computational technologies that are the main-stream in the future electronic structure calculations.

There are also many-body wave function methods that are very close to the quantum chemical methods. One example of these is the full configuration interaction method of chemistry that is termed the exact diagonalization (ED) in physics. The activities in porting the quantum chemistry approaches to GPU are reviewed in [89], and we try to remain on the physics side of this unclear borderline. We omit, for example, works on the coupled cluster method on GPU [90]. Furthermore, also quantum mechanical transport problems are not discussed here [91].

Lattice models [92,93] are important for providing a general understanding of many central physical concepts like magnetism. Furthermore, realistic materials can be cast to a lattice model [94]. Few-site models can be calculated exactly using the ED method. The ED method turns out to be very efficient on the GPU [95]. In the simplest form of ED, the first steps is to construct the many-body basis and the Hamiltonian matrix in it. Then follows the most time-consuming part, namely the actual diagonalization of the Hamiltonian matrix. In many cases, one is mainly interested in the lowest eigenstate and possibly a few of the lowest excited states. For these, the Lanczos algorithm turns out to be very suitable [95]. The basic idea of the Lanczos scheme is to map the huge but sparse Hamiltonian matrix to a smaller and tridiagonal form in the so called Krylov space that is defined by the spanning vectors, obtained from the starting vector  $|x_0\rangle$  by acting with the Hamiltonian as  $H^m|x_0\rangle$ . Now, as the GPU is very powerful for the matrix-vector product, it is not surprising that speedups of up to 100 can be found compared to CPUs [95].

## 6 Outlook

The GPU has made a definite entry into the world of computational physics: the appearance of studies that mention the GPU as only side note are a clear indicator of this, as preliminary studies using emerging technologies will always be done, but the true litmus test of a new technology is whether studies emerge where the new technology is actually used to advance science.

From the point of view of high performance computing in computational physics the biggest challenge facing GPUs at the moment is scaling: in the strong scaling case as many levels of parallelism as possible inherent in the problem should be exploited in order to reach for best performance with small local subsystems. The basic variables of the model are almost always vectors of some sort, making them an ideal candidate for SIMD-type of parallelism and this is often achieved with CPUs with a simple compiler-flag, which instructs the compiler to look for opportunities to combine independent instructions into vector-operations.

Furthermore, large and therefore interesting problems from a HPC point of view are typically composed of a large number of similar variables, be it particles, field values, cells or just entries in an array of numbers, which hints at another, higher level of parallelism of the problem that traditionally has been exploited using MPI, but is a prime candidate for a data-parallel algorithm. Also,

algorithmic changes may be necessary to reach the best possible performance: it may very well be that the best algorithm for CPUs is no longer the best one for GPUs; A classic example could be the question whether to use lookup tables of certain variables, or recompute them on-the-fly – typically on the GPU the FLOPS are cheap making the recomputation an attractive choice whereas the large caches of the CPU may make the lookup table a better option.

On the other hand, MPI-communication latencies should be minimized and bandwidth increased to accommodate the faster local solve to help with both weak and strong scaling. As far as we know there are very few, if any, groups taking advantage of GPUDirect v.2 for NVIDIA GPUs [96], which allows direct GPU-to-GPU MPI communications, reducing overhead and CPU synchronization needs. Even GPUDirect v.1 helps, as then one can share the *pinned memory* buffers between Infiniband and GPU cards, removing the need to do extra local copies of data. The MPI implementations should also be scaled to fit the needs of the GPUs connected to the node.

Another, perhaps an even greater challenge, facing GPUs and similar systems is the ecosystem: Currently a large portion of the developers and system administrators like to think of GPUs and similar solutions as *accelerators* – an accelerator is a component, which is attached to the main processor and used to speed up certain portions of the code, but as these “accelerators” become more and more agile with wider support for standard algorithms, the term becomes more and more irrelevant as a major part of the entire computation can be done on the “accelerator” and the original “brains” of the machine, the CPU, is mainly left there to take care of administrative functions, such as disk IO and common OS services.

As single threaded performance has reached a local limit, all types of processors are seeking more performance out of parallelism: more cores are added and vector units are broadened. This trend, fueled by the fact that transistor feature sizes keep on shrinking, hints at some type of convergence in the near future, but exactly what it will look like is anyone’s best guess. At least in computational physics, it has been shown already that the scientists are willing to take extra effort in porting their code to take advantage of massively parallel architectures, which should allow them to do the same work with less energy and do more science with the resources allocated to them.

The initial programming effort does raise a concern for productivity: How much time and effort is one willing to spend to gain a certain amount of added performance? Obviously, the answer depends on the problem itself, but perhaps even more on the assumed direction of the industry – a wrong choice may result in wasted effort if the chosen solution simply does not exist in five years time and this may depend strongly on consumer behavior, since a large part of the development costs of these machines are actually subsidized by the development of the consumer variants of the products. Designing a processor only for the HPC market is too expensive and a successful product will need a sister or at least a cousin in the consumer market. Which brings us back to DOOM: it may



very well be that the technology developed for the gamers of today, will be the programming platform for the scientists of tomorrow.

## Acknowledgements

We would like to thank Kari Rummukainen, Adam Foster, Risto Nieminen, Martti Puska, and Ville Havu for all their support. Topi Siro acknowledges the financial support from the Finnish Doctoral Programme in Computational Sciences FICS. This research has been partly supported by the Academy of Finland through its Centres of Excellence Program (Project No. 251748).

## References

1. Macedonia, M.: The GPU enters computing's mainstream. *Computer* **36**(10) (2003) 106–108
2. NVIDIA Corporation: NVIDIA CUDA C programming guide (2012) Version 4.2.
3. McCammon, J.A., Gelin, B.R., Karplus, M.: Dynamics of folded proteins. *Nature* **267**(5612) (1977) 585–590
4. Tembre, B.L., Cammon, J.M.: Ligand-receptor interactions. *Computers & Chemistry* **8**(4) (1984) 281 – 283
5. Gao, J., Kuczera, K., Tidor, B., Karplus, M.: Hidden thermodynamics of mutant proteins: a molecular dynamics analysis. *Science* **244**(4908) (1989) 1069–1072
6. Samish, I., MacDermaid, C.M., Perez-Aguilar, J.M., Saven, J.G.: Theoretical and computational protein design. *Annual Review of Physical Chemistry* **62**(1) (2011) 129–149
7. Berkowitz, M.L., Kindt, J.T. In: *Molecular Detailed Simulations of Lipid Bilayers*. John Wiley & Sons, Inc. (2010) 253–286
8. Lyubartsev, A.P., Rabinovich, A.L.: Recent development in computer simulations of lipid bilayers. *Soft Matter* **7** (2011) 25–39
9. Springel, V., White, S.D.M., Jenkins, A., Frenk, C.S., Yoshida, N., Gao, L., Navarro, J., Thacker, R., Croton, D., Helly, J., Peacock, J.A., Cole, S., Thomas, P., Couchman, H., Evrard, A., Colberg, J., Pearce, F.: Simulations of the formation, evolution and clustering of galaxies and quasars. *Nature* **435**(7042) (2005) 629–636
10. Bakker, A., Gilmer, G., Grabow, M., Thompson, K.: A special purpose computer for molecular dynamics calculations. *Journal of Computational Physics* **90**(2) (1990) 313 – 335
11. Susukita, R., Ebisuzaki, T., Elmegreen, B.G., Furusawa, H., Kato, K., Kawai, A., Kobayashi, Y., Koishi, T., McNiven, G.D., Narumi, T., Yasuoka, K.: Hardware accelerator for molecular dynamics: Mdgrape-2. *Computer Physics Communications* **155**(2) (2003) 115 – 131
12. Nishikawa, Y., Koibuchi, M., Yoshimi, M., Shitara, A., Miura, K., Amano, H.: Performance analysis of clearspeed's csx600 interconnects. In: *Parallel and Distributed Processing with Applications, 2009 IEEE International Symposium on*. (2009) 203–210
13. Harada, T.: GPU Gems 3: chapter 29. Volume 3. Addison-Wesley Professional (2007)
14. Le Grand, S.: GPU Gems 3: chapter 32. Volume 3. Addison-Wesley Professional (2007)

15. Chinchilla, F., Gamblin, T., Sommervoll, M., Prins, J.: Parallel N-body simulation using GPUs. Technical report, University of North Carolina (2004)
16. Brodtkorb, A.R., Hagen, T.R., Sætra, M.L.: Graphics processing unit (GPU) programming strategies and trends in GPU computing. *Journal of Parallel and Distributed Computing* (0) (2012) –
17. Stone, J.E., Hardy, D.J., Ufimtsev, I.S., Schulten, K.: GPU-accelerated molecular modeling coming of age. *Journal of Molecular Graphics and Modelling* **29**(2) (2010) 116 – 125
18. Nyland, L., Harris, M., Prins, J.: GPU Gems 3. chapter 31: Fast N-Body Simulation with CUDA. Volume 3. Addison-Wesley Professional (2007)
19. Allen, M.P., Tildesley, D.J.: *Computer Simulation of Liquids*. Oxford: Clarendon (2002)
20. Kipfer, P., Segal, M., Westermann, R.: Uberflow: a GPU-based particle engine. In: *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware*. HWWS '04, New York, NY, USA, ACM (2004) 115–122
21. Satish, N., Harris, M., Garland, M.: Designing efficient sorting algorithms for manycore GPUs. Technical report, NVIDIA (2008)
22. Ha, L., Krüger, J., Silva, C.T.: Fast four-way parallel radix sorting on GPUs. *Computer Graphics Forum* **28**(8) (2009) 2368–2378
23. Anderson, J.A., Lorenz, C.D., Travesset, A.: General purpose molecular dynamics simulations fully implemented on graphics processing units. *Journal of Computational Physics* **227**(10) (2008) 5342 – 5359
24. Moreland, K., Angel, E.: The FFT on a GPU. In: *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware*. HWWS '03, Aire-la-Ville, Switzerland, Switzerland, Eurographics Association (2003) 112–119
25. Govindaraju, N.K., Manocha, D.: Cache-efficient numerical algorithms using graphics hardware. Technical report, The University of North Carolina (2007)
26. Gu, L., Li, X., Siegel, J.: An empirically tuned 2D and 3D FFT library on CUDA GPU. In: *Proceedings of the 24th ACM International Conference on Supercomputing*. ICS '10, New York, NY, USA, ACM (2010) 305–314
27. Chen, Y., Cui, X., Mei, H.: Large-scale FFT on GPU clusters. In: *Proceedings of the 24th ACM International Conference on Supercomputing*. ICS '10, New York, NY, USA, ACM (2010) 315–324
28. Ahmed, M., Haridy, O.: A comparative benchmarking of the FFT on Fermi and Evergreen GPUs. In: *Performance Analysis of Systems and Software (ISPASS)*, 2011 IEEE International Symposium on. (2011) 127 –128
29. Skeel, R.D., Tezcan, I., Hardy, D.J.: Multiple grid methods for classical molecular dynamics. *Journal of Computational Chemistry* **23**(6) (2002) 673–684
30. Hardy, D.J., Stone, J.E., Schulten, K.: Multilevel summation of electrostatic potentials using graphics processing units. *Parallel Computing* **35**(3) (2009) 164 – 177
31. Goodnight, N., Woolley, C., Lewin, G., Luebke, D., Humphreys, G.: A multigrid solver for boundary value problems using programmable graphics hardware. In: *ACM SIGGRAPH 2005 Courses*. SIGGRAPH '05, New York, NY, USA, ACM (2005)
32. McAdams, A., Sifakis, E., Teran, J.: A parallel multigrid poisson solver for fluids simulation on large grids. In: *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. SCA '10, Aire-la-Ville, Switzerland, Switzerland, Eurographics Association (2010) 65–74

33. Rensselaer Polytechnic Institute. Image Processing Laboratory, Meagher, D.: Octree Encoding: a New Technique for the Representation, Manipulation and Display of Arbitrary 3-D Objects by Computer. (1980)
34. Lefebvre, S., Hornus, S., Neyret, F.: GPU Gems 2. chapter 37: Octree Textures on the GPU. Volume 2. Addison-Wesley Professional (2005)
35. Belleman, R.G., Bédorf, J., Zwart, S.F.P.: High performance direct gravitational N-body simulations on graphics processing units ii: An implementation in CUDA. *New Astronomy* **13**(2) (2008) 103 – 112
36. Hamada, T., Narumi, T., Yokota, R., Yasuoka, K., Nitadori, K., Taiji, M.: 42 tflops hierarchical N-body simulations on GPUs with applications in both astrophysics and turbulence. In: Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis. SC '09, New York, NY, USA, ACM (2009) 62:1–62:12
37. Rokhlin, V.: Rapid solution of integral equations of classical potential theory. *Journal of Computational Physics* **60**(2) (1985) 187 – 207
38. Greengard, L., Rokhlin, V.: A fast algorithm for particle simulations. *Journal of Computational Physics* **73**(2) (1987) 325 – 348
39. Gumerov, N.A., Duraiswami, R.: Fast multipole methods on graphics processors. *Journal of Computational Physics* **227**(18) (2008) 8290 – 8313
40. Darve, E., Cecka, C., Takahashi, T.: The fast multipole method on parallel clusters, multicore processors, and graphics processing units. *Comptes Rendus Mécanique* **339**(2-3) (2011) 185 – 193
41. Takahashi, T., Cecka, C., Fong, W., Darve, E.: Optimizing the multipole-to-local operator in the fast multipole method for graphical processing units. *International Journal for Numerical Methods in Engineering* **89**(1) (2012) 105–133
42. Yokota, R., Barba, L., Narumi, T., Yasuoka, K.: Petascale turbulence simulation using a highly parallel fast multipole method on GPUs. *Computer Physics Communications* (0) (2012) –
43. Yasuda, K.: Accelerating density functional calculations with graphics processing unit. *Journal of Chemical Theory and Computation* **4**(8) (2008) 1230–1236
44. Yasuda, K.: Two-electron integral evaluation on the graphics processor unit. *Journal of Computational Chemistry* **29**(3) (2008) 334–342
45. Genovese, L., Ospici, M., Deutsch, T., Méhaut, J.F., Neelov, A., Goedecker, S.: Density functional theory calculation on many-cores hybrid central processing unit-graphic processing unit architectures. *The Journal of Chemical Physics* **131**(3) (2009) 034103
46. Genovese, L., Neelov, A., Goedecker, S., Deutsch, T., Ghasemi, S.A., Willand, A., Caliste, D., Zilberberg, O., Rayson, M., Bergman, A., Schneider, R.: Daubechies wavelets as a basis set for density functional pseudopotential calculations. *The Journal of Chemical Physics* **129**(1) (2008) 014109
47. Kresse, G., Furthmüller, J.: Efficient iterative schemes for *ab initio* total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **54** (Oct 1996) 11169–11186
48. Maintz, S., Eck, B., Dronskowski, R.: Speeding up plane-wave electronic-structure calculations using graphics-processing units. *Computer Physics Communications* **182**(7) (2011) 1421 – 1427
49. Hacene, M., Anciaux-Sedrakian, A., Rozanska, X., Klahr, D., Guignon, T., Fleurat-Lessard, P.: Accelerating VASP electronic structure calculations using graphic processing units. *Journal of Computational Chemistry* (2012) n/a–n/a
50. Giannozzi, P., Baroni, S., Bonini, N., Calandra, M., Car, R., Cavazzoni, C., Ceresoli, D., Chiarotti, G.L., Cococcioni, M., Dabo, I., Corso, A.D., de Gironcoli,

- S., Fabris, S., Fratesi, G., Gebauer, R., Gerstmann, U., Gougoussis, C., Kokalj, A., Lazzeri, M., Martin-Samos, L., Marzari, N., Mauri, F., Mazzarello, R., Paolini, S., Pasquarello, A., Paulatto, L., Sbraccia, C., Scandolo, S., Sclauzero, G., Seitsonen, A.P., Smogunov, A., Umari, P., Wentzcovitch, R.M.: Quantum espresso: a modular and open-source software project for quantum simulations of materials. *Journal of Physics: Condensed Matter* **21**(39) (2009) 395502
51. Girotto, I., Varini, N., Spiga, F., Cavazzoni, C., Ceresoli, D., Martin-Samos, L., Gorni, T.: Enabling of Quantum-ESPRESSO to petascale scientific challenges. In: PRACE Whitepapers. PRACE (2012)
  52. Spiga, F., Girotto, I.: phiGEMM: A CPU-GPU library for porting Quantum ESPRESSO on hybrid systems. In: Parallel, Distributed and Network-Based Processing (PDP), 2012 20th Euromicro International Conference on. (feb. 2012) 368–375
  53. Wang, L., Wu, Y., Jia, W., Gao, W., Chi, X., Wang, L.W.: Large scale plane wave pseudopotential density functional theory calculations on GPU clusters. In: Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis. SC '11, New York, NY, USA, ACM (2011) 71:1–71:10
  54. Enkovaara, J., Rostgaard, C., Mortensen, J.J., Chen, J., Dulak, M., Ferrighi, L., Gavnholt, J., Glinsvad, C., Haikola, V., Hansen, H.A., Kristoffersen, H.H., Kuisma, M., Larsen, A.H., Lehtovaara, L., Ljungberg, M., Lopez-Acevedo, O., Moses, P.G., Ojanen, J., Olsen, T., Petzold, V., Romero, N.A., Stausholm-Mller, J., Strange, M., Tritsarlis, G.A., Vanin, M., Walter, M., Hammer, B., Hkkinen, H., Madsen, G.K.H., Nieminen, R.M., Nrskov, J.K., Puska, M., Rantala, T.T., Schitz, J., Thygesen, K.S., Jacobsen, K.W.: Electronic structure calculations with GPAW: a real-space implementation of the projector augmented-wave method. *Journal of Physics: Condensed Matter* **22**(25) (2010) 253202
  55. Peskin, M.E., Schroeder, D.V.: An Introduction to Quantum Field Theory. Westview Press (1995)
  56. Crewther, R.J.: Introduction to Quantum Field Theory. ArXiv High Energy Physics - Theory e-prints (1995)
  57. Fodor, Z., Hoelbling, C.: Light Hadron Masses from lattice QCD. *Reviews of Modern Physics* **84** (2012) 449–495
  58. Gckeler, M., Hgler, P., Horsley, R., Pleiter, D., Rakow, P.E.L., Schfer, A., Schierholz, G., Zanotti, J.M.: Generalized parton distributions and structure functions from full lattice QCD. *Nuclear Physics B Proceedings Supplements* **140** (2005) 399–404
  59. Alexandrou, C., Brinet, M., Carbonell, J., Constantinou, M., Guichon, P., et al.: Nucleon form factors and moments of parton distributions in twisted mass lattice QCD. In: Proceedings of The XXIst International Europhysics Conference on High Energy Physics. July 21-27 2011. Grenoble, Rhones Alpes France. Volume EPS-HEP2011. (2011) 308
  60. McNeile, C., Davies, C.T.H., Follana, E., Hornbostel, K., Lepage, G.P.: High-precision  $f_{B_s}$  and heavy quark effective theory from relativistic lattice QCD. *Physical Review D* **85** (2012) 031503
  61. Rummukainen, K.: QCD-like technicolor on the lattice. In Llanes-Estrada, F.J., Peláez, J.R., eds.: American Institute of Physics Conference Series. Volume 1343 of American Institute of Physics Conference Series. (2011) 51–56
  62. Petreczky, P.: Progress in finite temperature lattice QCD. *Journal of Physics G: Nuclear and Particle Physics* **35**(4) (2008) 044033

63. Petreczky, P.: Recent progress in lattice QCD at finite temperature. ArXiv e-prints (2009)
64. Montvay, I., Münster, G.: Quantum Fields on a Lattice. Cambridge Monographs on Mathematical Physics. Cambridge University Press, The Edinburgh Building, Cambridge CB2 2RU, United Kingdom (1994)
65. Rothe, H.J.: Lattice Gauge Theories: An Introduction. 3 edn. World Scientific Publishing Company, Hackensack, N.J. (2005)
66. Gupta, R.: Introduction to Lattice QCD. ArXiv High Energy Physics - Lattice e-prints (1998)
67. Egri, G., Fodor, Z., Hoelbling, C., Katz, S., Nogradi, D., Szabo, K.: Lattice QCD as a video game. Computer Physics Communications **177** (2007) 631–639
68. Schröck, M., Vogt, H.: Gauge fixing using overrelaxation and simulated annealing on GPUs. ArXiv e-prints (2012)
69. Babich, R., Clark, M.A., Joó, B., Shi, G., Brower, R.C., Gottlieb, S.: Scaling lattice QCD beyond 100 GPUs. In: Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis. SC '11, New York, NY, USA, ACM (2011) 70:1–70:11
70. Hasenbusch, M., Jansen, K.: Speeding up the HMC: QCD with clover-improved wilson fermions. Nuclear Physics B Proceedings Supplements **119** (2003) 982–984
71. Osaki, Y., Ishikawa, K.I.: Domain Decomposition method on GPU cluster. In: Proceedings of The XXVIII International Symposium on Lattice Field Theory. June 14-19, 2010. Villasimius, Sardinia Italy. (2010)
72. Bonati, C., Cossu, G., D'Elia, M., Incardona, P.: QCD simulations with staggered fermions on GPUs. Computer Physics Communications **183** (2012) 853–863
73. Winter, F.: Accelerating QDP++ using GPUs. In: Proceedings of the XXIX International Symposium on Lattice Field Theory (Lattice 2011). July 10-16, 2011. Squaw Valley, Lake Tahoe, California. (2011)
74. Walk, B., Wittig, H., Dranschnikow, E., Schomer, E.: Implementation of the Neuberger overlap operator in GPUs. In: Proceedings of The XXVIII International Symposium on Lattice Field Theory. June 14-19, 2010. Villasimius, Sardinia Italy. (2010)
75. Alexandru, A., Lujan, M., Pelissier, C., Gamari, B., Lee, F.X.: Efficient implementation of the overlap operator on multi-GPUs. ArXiv e-prints (2011)
76. Cardoso, N., Bicudo, P.: SU (2) lattice gauge theory simulations on Fermi GPUs. Journal of Computational Physics **230** (2011) 3998–4010
77. Cardoso, N., Bicudo, P.: Generating SU(Nc) pure gauge lattice QCD configurations on GPUs with CUDA. ArXiv e-prints (2011)
78. Amado, A., Cardoso, N., Cardoso, M., Bicudo, P.: Study of compact U(1) flux tubes in 3+1 dimensions in lattice gauge theory using GPU's. ArXiv e-prints (2012)
79. Bordag, M., Demchik, V., Gulov, A., Skalozub, V.: The type of the phase transition and coupling values in  $\lambda\phi^4$  model. International Journal of Modern Physics A **27** (2012) 50116
80. Chiu, T.W., Hsieh, T.H., Mao, Y.Y.: Pseudoscalar Meson in Two Flavors QCD with the Optimal Domain-Wall Fermion. Physics Letters B **B717** (2012) 420
81. Munshi, A.: The OpenCL specification (2011) Version 1.2.
82. Bach, M., Lindenstruth, V., Philipsen, O., Pinke, C.: Lattice QCD based on OpenCL. ArXiv e-prints (2012)
83. IBM Systems and Technology: IBM System Blue Gene/Q – Data Sheet (2011)
84. Foulkes, W.M.C., Mitas, L., Needs, R.J., Rajagopal, G.: Quantum monte carlo simulations of solids. Reviews of Modern Physics **73** (2001) 33–83

85. Harju, A., Barbiellini, B., Siljamäki, S., Nieminen, R., Ortiz, G.: Stochastic gradient approximation: An efficient method to optimize many-body wave functions. *Physical Review Letters* **79**(7) (1997) 1173–1177
86. Harju, A.: Variational Monte Carlo for interacting electrons in quantum dots. *Journal of Low Temperature Physics* **140**(3-4) (2005) 181–210
87. Anderson, A.G., III, W.A.G., Schröder, P.: Quantum Monte Carlo on graphical processing units. *Computer Physics Communications* **177**(3) (2007) 298 – 306
88. Esler, K., Kim, J., Ceperley, D., Shulenburger, L.: Accelerating quantum Monte Carlo simulations of real materials on GPU clusters. *Computing in Science and Engineering* **14**(1) (2012) 40–51
89. Wölfe, A.W.G., Walker, R.C.: Chapter 2 - quantum chemistry on graphics processing units. In Wheeler, R.A., ed.: *Annual Reports in Computational Chemistry*. Volume 6 of *Annual Reports in Computational Chemistry*. Elsevier (2010) 21 – 35
90. DePrince, A., Hammond, J.: Quantum chemical many-body theory on heterogeneous nodes. In: *Application Accelerators in High-Performance Computing (SAAHPC)*, 2011 Symposium on. (2011) 131 –140
91. Ihnatsenka, S.: Computation of electron quantum transport in graphene nanoribbons using GPU. *Computer Physics Communications* **183**(3) (2012) 543 – 546
92. Hubbard, J.: Electron correlations in narrow energy bands. *Proc. R. Soc. Lond. Ser. A* **276**(1364) (1963) 238–257
93. Gutzwiller, M.C.: Effect of correlation on the ferromagnetism of transition metals. *Phys. Rev. Lett.* **10** (1963) 159–162
94. Meredith, J.S., Alvarez, G., Maier, T.A., Schulthess, T.C., Vetter, J.S.: Accuracy and performance of graphics processors: A quantum Monte Carlo application case study. *Parallel Computing* **35**(3) (2009) 151 – 163
95. Siro, T., Harju, A.: Exact diagonalization of the Hubbard model on graphics processing units. *Computer Physics Communications* **183**(9) (2012) 1884 – 1889
96. NVIDIA Corporation: NVIDIA GPUDirect<sup>TM</sup> Technology (2012)